

LSBR Scientific Advisory Meeting: Computing

NSLS-II, BNL, Upton, NY, March 28 2017




EIGER 9M @ AMX and 16M @ FMX



Eiger9M
18 modules
10 M pixels
238 Hz


Eiger16M
32 modules
18 M pixels
133 Hz

Eiger 4M ROI
8 modules
4.5 M pixels
750 Hz



233x245	311x328	mm ²
238 (750)	133 (750)	Hz
10.2	18.1	M pixel
< 5	< 8	M byte (actual)
200-1700	200-1700	MB/s

2-30 TB / day / beamline



Challenges : data storage, transfer, processing and backup.
Requires advanced software and computing cluster.

AMX / FMX: Current status (275 mA)

AMX : from 5 to 18 keV

Flux > 3×10^{12} ph.s⁻¹

Beam ~ 8 x 7 μm^2

Manual sample mounting (up to 10/H)

Typical data collection:

0.1 deg @ 100 Hz

Data collection : from <1s to 15 secs

Default screening:

10 % transmission & 10 ms/frame

Av data set: 1400 frames

500 frames / data.h5

Compression range: 4.2-25

Aver compression:10

FMX : from 5 to 25 keV

Flux > 2×10^{12} ph.s⁻¹

Beam ~ 10 x 6 μm^2 (as of last week)

Automated sample mounting (up to 20/H)

Typical data collection:

0.1 deg @ 50 Hz

Data collection : from <1s to 15 secs

Default screening:

<10 % transmission & 50 ms/frame

Av data set: 1800 frames

200 frames / data.h5

Compression range: 4.5-32

Aver compression:11.6

Operate 2 shifts / day, at the moment.

~ 500 GB data per 2 shifts on AMX (at the moment)

Computing / Network Infrastructures

9M@AMX

EIGER
Detector

16M@FMX

x controls
Area detector
IOC server
(GPFS / 10 Gb)

3 x Data
Collection
Processing
Workstations
(GPFS / 10 Gb)

FMX & AMX

40 Gb/s

Compute Nodes
(LSBR)



208 cores
4+4 nodes

GPFS Storage
(NSLS-II)



860 TB
IB

GPFS Buffer
(LSBR)



20 TB
IB

NSLS-II Computing Facility

10 then 40 Gb/s

> 2000 KNL cores; > 1 PB GPFS



BNL Institutional Cluster

Data collection using LSDC.

Data visualization using Albula or ADSC.

Manual data processing using HKL2000 or other
GUIs (ccp4i, phenix, imosflm ...), index/Best.

Data transfer onto USB drives.

Direct access to compute nodes over terminal.

Rapid Data Analysis and Reduction on dedicated
compute nodes and fast buffer

Pipelines and packages in use: dials spot finders,
dozor, XDS, dials, xia2, fast_dp, fast_mr,
dimple, fast_ep, fast_ep_weak, fast_ep_NSLS-
II, KAMO, and more.

Network backup using dedicated node.

Post processing, Re-processing, Remote
processing, Multiple Crystals data reduction
and Hierarchical Cluster Analysis.

To world (1 Gb/s (internet) to 100 Gb/s (EsNet))

Rastering: goal is near real time analysis

1 master/data file per row.


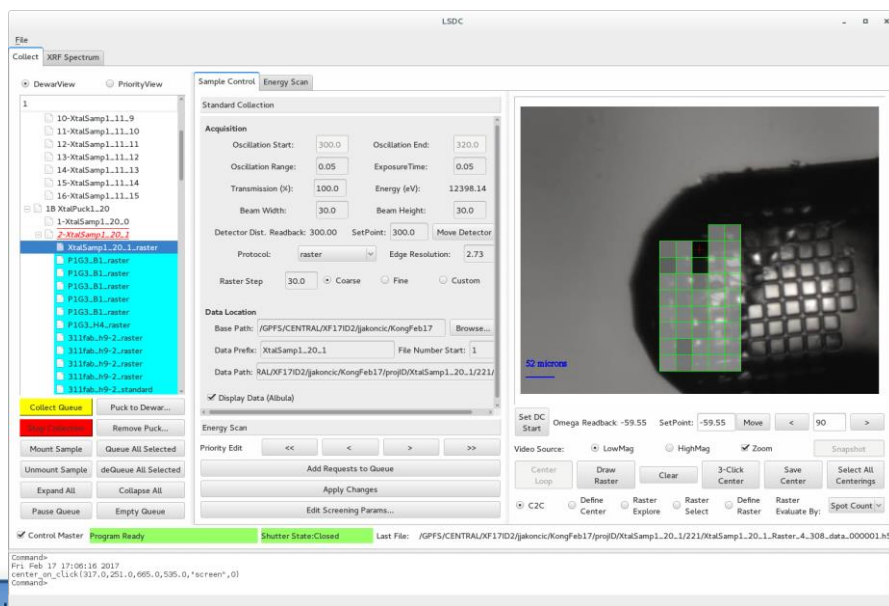
Frames converted (eiger2cbf) and processed with dials.find_spots client / server (ping/pong 2 nodes)

Data located on GPFS

Current performances: processing at ~ 15 FPS on one node
(need to be at ~ 200 E9M FPS, ~ 100 E16M FPS and ~ 500 E4M FPS)

Mixed results when ice rings are present: optimizing dials code and or using dozor (ESRF/EMBL)

>>>> optimizing codes for faster results (ROI, HDF5 read, KNL, OMP




fine

Data Reduction: goal is dataset reduction under 2 min

Fast dp: on 1 dedicated node (more to be used soon) @ 2-10 FPS.

data set of 900 frames (lysozyme test). (all data on GPFS)

9M	Total time (s)	FPS
FastDP-002	86 (79 RAM)	10 (11)
FastDP-010	83 (50 RAM)	11 (18)
XDS-002	61 (55 RAM)	15 (16)
XDS-010	72 (43 RAM)	12 (21)



XDS: jobs = all, 72 procs, 16 jobs)

Average performances on GPFS
(depending on load: IO dependant)

CPU04, 005: 36 [cores@2.3](#) GHz (10 Gb)

CPU09, 010: 44 [cores@2.2](#) GHz (40 Gb)

CPU02, 003, 007, 008: 12 [cores@3.4](#) GHz (10 Gb)

NetApp EF560 All-Flash Array



20 x 1.6 TB all flash

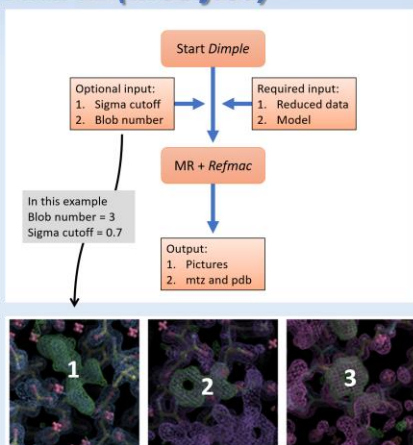
6 GB/s (50/50)

2 IB controllers 2 ports each

Procured/Delivered/ To be Tested

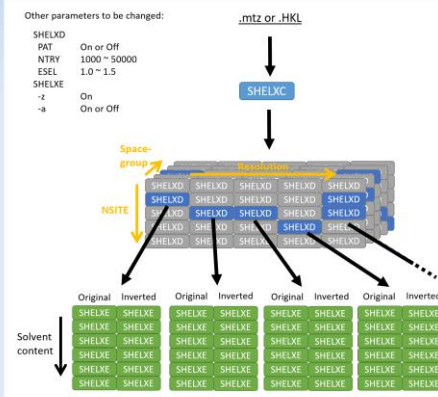
Analysis Pipelines

DIMPLE (modified)

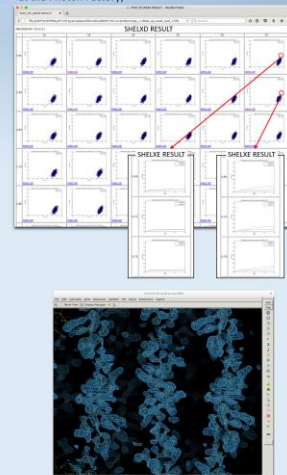


FAST_EP_WEAK

Fast_ep_weak is a modified version of Fast_ep, which is originally developed at Diamond Light Source. Fast_ep_weak is designed to find a structure solution from weak anomalous signals by finer grid searches with spacegroup, nsite and resolution cut off.



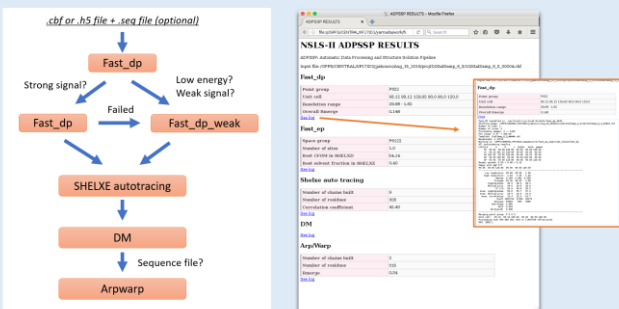
P-SAD on a DNA crystal (Data sets were collected at BL1A at the Photon Factory)



A. Soares, Y. Yamada (PF)
G.Winter (DLS)., H.Bernstein (RIT)


FAST_EP_NSL2 (NSLS-II ADPSSP)

Fast_ep_nsl2 is a single crystal SAD phasing pipeline from a diffraction data set to a model with fast_dp, fast_ep, fast_ep_weak and other crystallographic software.



Thermolysin Zn-SAD at AMX at 11 keV


fast_dp (data reduction)	< 350 s	fast_ep	< 200 s
Source group	PS22	OP1	0.049
Unit-cell parameters	a=b=93.13 c=129.92	d(I)/d(O)	0.888
Mosaicity spread (°)	0.05 < 0.08 < 0.12	Best space group	P6 ₃ /2
Resolution (Å)	29.68-1.51 (1.55-1.51)	SHELXE autotracing	6 m
R _{merge}	0.121 (0.627)	CC (%)	42.42
/Sigma	21.20 (3.30)	Chains	8
Completeness (%)	99.7 (96.6)	Residues	306
Multiplicity	25.1 (12.6)	arpWarp (2 building cycles)	< 130 s
CC(I)	99.9 (88.1)	Residues / water	314 (out of 316) / 0
Anom. Completeness	99.7 (96.0)	R _{merge} / R _{free}	0.37 / 0.40
Resolution (Å)	29.68-1.51 (1.55-1.51)	arpWarp (5 building cycles)	< 600 s
R _{merge}	0.121 (0.627)	Residues / water	315 / 330
/Sigma	21.20 (3.30)	R _{merge} / R _{free}	0.22 / 0.25



All processing performed on a single node: total time < 15 mins

KAMO

Multiple small wedge data collection at AMX



Energy	8
Flux	5.3 x 10 ¹¹ ph/sec (85 % attenuation)
Total oscillation	5 deg./spot
Osc. width	0.2 deg./frame
Exp. time	0.04 sec./frame (P6M)

(multi.collect on LSDC)
Raster vs raster + threshold

Collaborations

H. Bernstein (RIT), Y. Yamada (PF), S. Popov (ESRF), G. Bourenkov (EMBL), G. Winter (DLS), N. Sauter (ALS).

Spot finder (dozor and dials): read HDF5 using MP


Indexing (labelit.index) on eiger data producing best files:

Best Strategy on eiger data AND based on data previously processed

> multiple crystals strategy

Data reduction (XDS): parallel colspot, read HDF5 (ram mapping), KNL implementation

HDRMX : sharing info slurm, fast dp,



HDRMX meeting in MAX-IV, Lund (not in picture: Y. Yamada, S. Popov and N. Sauter)

To keep up with the data:

40 Gb/s (infiniband) on all nodes (fast IO) ;

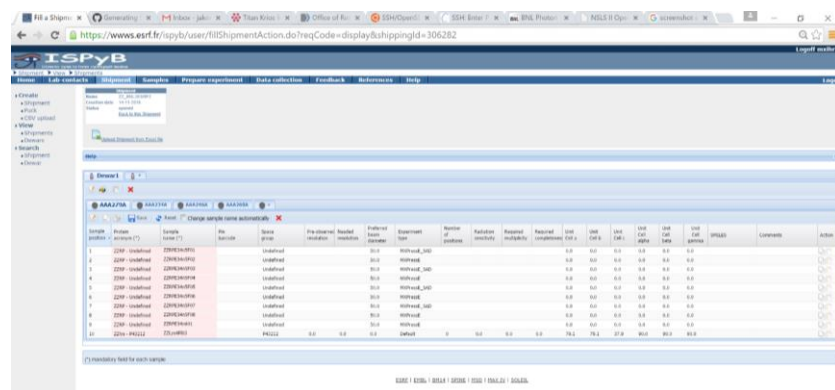
2nd buffer disk (as needed)

Additional compute nodes: > 4 nodes (160 physical cores)

Full slurm deployment (cluster workload management with GPL license)

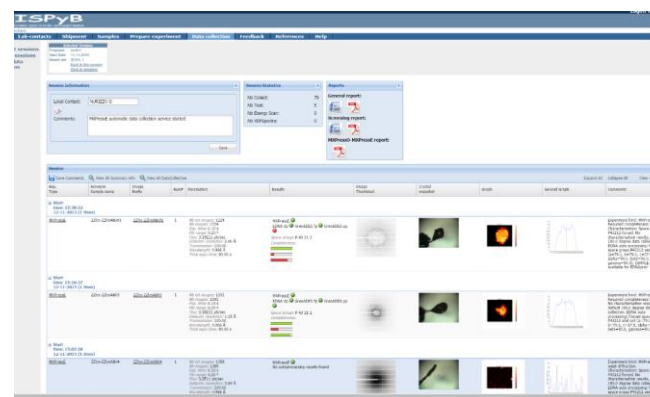
Database (such as IspyB) is an absolute requirement (Bioinformatics (2011) 27 (22): 3186-3192.);

Users can't keep up with data processing and data quality assessment



The screenshot shows the IspyB web interface with a table of sample data. The table has columns for Sample ID, Protein, Sample, Site, Date, Pre-processed, Sample, Experiment, Number of proteins, Radiation dose, Radiation dose, Radiation dose, and various quality metrics. The table is filtered for 'Sample ID' and shows 10 rows of data.

Sample ID	Protein	Sample	Site	Date	Pre-processed	Sample	Experiment	Number of proteins	Radiation dose	Radiation dose	Radiation dose	Quality metrics
1	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10
2	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10
3	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10
4	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10
5	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10
6	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10
7	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10
8	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10
9	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10
10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10	2019-10-10-10-10-10



Implementation of additional pipelines, specifically for multiple crystals

Using KNLs: spot finder, dozor, xds, dials ? We are testing on Intel cluster: ACA workshop with Intel

Conclusions & Acknowledgements

Data retention: at the moment not an issue (GPFS capacity)

Remote access will be implemented in Summer (NoMachine)

With beams 10 microns or less, we see more head scratching.

Data collection from multiple micro-crystals more and more used.

MR and SAD pipelines are executed manually at the moment (user input required)

Not all users are prepared to handle the amount of data and reprocessing data :
database a must and general cluster with remote access would help.

Learning as we go and implement new features (automation, remote, Multiple crystals data collections, in-situ ...

LIX Status: GPFS in use, work in progress with DAMA group.

J. Skinner, B. Martins, R. Petkus, M. Fuchs, A. Soares, E. Lazo.
N. Sauter, G. Winter, S. Popov, G. Bourenkov.
W. Shi, H. Bernstein, Y. Yamada[#].



Save the date:

Open House (Proprietary): April 11

ACA Workshop: May26

NSLS-II: Expert User Training (BAG): 1st ½ June